COMPAQ
Integration Note

September 1999
0096-0699-A

Prepared by ECG Technology Communications Group

Compaq Computer Corporation

Contents

Clustering Overview 3
Clusters Defined 3
Causes of Downtime 3
Cost of Downtime 4
Benefits of Clustering 4
Summary 4
ProLiant Cluster for NetWare 4.2 Overview 5
Features and Benefits of the ProLiant Cluster for NetWare 4.2 5
Novell's High-Reliability Solution 5
The Failover Process 6
Hardware Requirements 8
Compaq ProLiant Servers 8
Compaq StorageWorks RAID Array 4000 Storage System 9
Cluster Interconnects 12
Software Components 13
Novell Software 13
Compaq Software 13
Planning Novell High Availability Server Installation 15
Verification of NHAS Installation 15
ProLiant Cluster for NetWare 4.2 Administration 18
Increasing Availability 18
IP and Applications Failover 18
Management 26
Technical FAQs 28

Deploying the ProLiant Cluster for NetWare 4.2

Abstract: Compaq is renowned in the Enterprise market with its industry-leading, highly available, fault-tolerant servers. Compaq servers and storage deliver on this reputation by offering the highest level of reliability and uptime with features such as:

- Industry standard PCI Hot Plug and internal hot-plug storage
- Redundant hot-plug fans, power supplies, and NICs
- Automatic Server Recovery-2 (ASR-2)
- ECC protected memory
- Integrated Management Display (IMD)
- Integrated Remote Console (IRC)
- Compaq SmartStart and Compaq Insigbt Manager

The Compaq ProLiant Cluster for NetWare 4.2 is designed to help customers reduce the risk and costs of downtime due to hardware or software failures. Benefits include:

- Automatic failover protection - In the event of a failure on one server, the remaining server maintains availability of network data and resources.
- Active servers - Users can now cluster two active NetWare 4.11 or NetWare 4.2 servers on the network, providing failover support and significant cost savings, since purchasing a dedicated server for failover is no longer necessary.
- Shared storage - Support for shared storage devices allows customers to reduce overall management costs by centralizing administration and back-up. Shared storage also allows more efficient use of storage space by eliminating the need to maintain duplicate files on different servers.
- No single point of failure – The new Compaq StorageWorks RAID Array 4000 storage system provides redundant fibre channel loops for No Single Point of Failure configurations.

The ProLiant Cluster for NetWare 4.2 provides the highest levels of system and application availability for customers with high system uptime requirements for NetWare 4.x environments. These levels are necessary to meet the requirements of demanding business-critical environments. This Integration Note details the configuration options and server requirements, provides insight into the administration, and describes some of the features available with the ProLiant Cluster for NetWare 4.2.

Downloaded from www.Manualslib.com manuals search engine

Notice

The information in this publication is subject to change without notice and is provided "AS IS" WITHOUT WARRANTY OF ANY KIND. THE ENTIRE RISK ARISING OUT OF THE USE OF THIS INFORMATION REMAINS WITH RECIPIENT. IN NO EVENT SHALL COMPAQ BE LIABLE FOR ANY DIRECT, CONSEQUENTIAL, INCIDENTAL, SPECIAL, PUNITIVE OR OTHER DAMAGES WHATSOEVER (INCLUDING WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION OR LOSS OF BUSINESS INFORMATION), EVEN IF COMPAQ HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

The limited warranties for Compaq products are exclusively set forth in the documentation accompanying such products. Nothing herein should be construed as constituting a further or additional warranty.

This publication does not constitute an endorsement of the product or products that were tested. The configuration or configurations tested or described may or may not be the only available solution. This test is not a determination or product quality or correctness, nor does it ensure compliance with any federal state or local requirements.

Product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.

Compaq, Contura, Deskpro, Fastart, Compaq Insight Manager, LTE, PageMarq, Systempro, Systempro/LT, ProLiant, TwinTray, ROMPaq, LicensePaq, QVision, SLT, ProLinea, SmartStart, NetFlex, DirectPlus, QuickFind, RemotePaq, BackPaq, TechPaq, SpeedPaq, QuickBack, PaqFax, Presario, SilentCool, CompaqCare (design), Aero, SmartStation, MiniStation, and PaqRap, registered United States Patent and Trademark Office.

Netelligent, Armada, Cruiser, Concerto, QuickChoice, ProSignia, Systempro/XL, NetI, LTE Elite, Vocalyst, PageMate, SoftPaq, FirstPaq, SolutionPaq, EasyPoint, EZ Help, MaxLight, MultiLock, QuickBlank, QuickLock, UltraView, Innovate logo, Wonder Tools logo in black/white and color, and Compaq PC Card Solution logo are trademarks and/or service marks of Compaq Computer Corporation.

Microsoft, Windows, Windows NT, Windows NT Server and Workstation, Microsoft SQL Server for Windows NT are trademarks and/or registered trademarks of Microsoft Corporation.

NetWare and Novell are registered trademarks and intraNetWare, NDS, Novell Directory Services, and High Availability Server (NHAS) are trademarks of Novell, Inc.

Pentium is a registered trademark of Intel Corporation.

Copyright ©1999 Compaq Computer Corporation. All rights reserved. Printed in the U.S.A.

Deploying the ProLiant Cluster for NetWare 4.2
Integration Note prepared by ECG Technology Communications Group

Second Edition (June 1999)
Document Number 0096-0699-A

Downloaded from www.Manualslib.com manuals search engine

Clustering Overview

The computer industry has been using a wide range of solutions to reduce system downtime for years. Many of these solutions have been both expensive and difficult to setup and maintain. Historically, only mission-critical applications, such as those controlling stock exchange trading floors and aerospace missions, were deemed important enough to justify proprietary clustering solutions.

However, as the presence of computer systems in day-to-day business operations has increased, the amount of acceptable downtime has decreased. Today, a new class of applications exists; business-critical applications key to business success but not significant enough to justify the high price tag of a proprietary clustering solution. More and more applications are becoming business-critical. Computer applications are becoming more widely deployed and their failure causes immediate lost revenue, decreased productivity, and, potentially, dissatisfied customers.

Poised to meet the increased demands, clustering technology is entering mainstream industry-standard computing. As a result, new clustering solutions, using industry-standard hardware and software, provide key clustering features at a lower price than proprietary clustering systems.

Clusters Defined

Traditionally, clustering is the integration of software and hardware that enables a set of loosely coupled servers and storage to operate as a single system, presenting a single entity to clients. As a cluster, the group of servers and storage offers a level of availability greatly exceeding the reliability of a standalone server. This translates into increased performance and greater data availability for end users.

The ProLiant Cluster for NetWare 4.2 extends the definition by providing the same level of reliability while maintaining distinct network entities for each server. By defining distinct server identities, tasks and users may be distributed easily between the distinct servers. With the focus for clustered systems being reliability, the ProLiant Cluster for NetWare 4.2 delivers again with automatic failover of all network resources and responsibilities in the event of a failure. Should one server in the cluster lose access to the shared storage and network, the other server assumes the responsibilities of the downed server.

Causes of Downtime

Computer downtime is the period of time that a computer system cannot meet the requests of its users. The following are the leading causes of downtime:

- *Planned service* - In a cluster, a single server can be taken offline, while the other server takes on the workload of the downed server. This allows planned service to occur with only minimal interruption to the overall system operation.
- *Software failures* - Because clustering provides a mechanism to automatically fail over processes when a discernible software failure occurs, the overall system operation can continue with only minimal interruption.
- *Hardware failures* - If it is determined that a hardware failure will result in system downtime, the cluster will fail over processes from one node to another, allowing the overall system operation to continue with little interruption.

Downloaded from www.Manualslib.com manuals search engine

Cost of Downtime

The driving force behind the high-availability systems has been the elimination of costs associated with downtime. The cost of downtime formula must include several factors. Among these are:

- Productivity loss
- IM costs
- Lost transactions
- Customer or end user dissatisfaction

Each factor is weighted differently, depending on the critical nature of each factor as it pertains to your business and to your specific application systems. For example, downtime during peak bours of a point-of-sale operation would have a much greater impact on customer satisfaction issues than downtime in an end-of-day email server backup operation. To understand the true cost of computer downtime in your business environment, you must examine the impact and contribution to increased costs of each factor as it applies to your business-critical applications.

Benefits of Clustering

With the causes enumerated and the costs understood, it is clear that the Novell High Availability Server can eliminate your concerns and ease your tightening networking budget through

- *Availability* - Simply defined, availability is the measure of how well a computer system can continuously deliver services to clients. This is dependent upon the system's ability to prevent or recover from failures or faults.
- *Scalability* - Clustering for scalability means increasing performance beyond that of a single computer node by adding more nodes to the cluster. Performance scalability across cluster nodes is difficult to achieve and requires not only scalable hardware, but also scalable software, for example, a parallel database.
- *Lower Total Cost of Ownership* - By reducing downtime, you can eliminate both the high costs of administrative support and, at the same time, the lost revenue caused by downtime and customer dissatisfaction.

Summary

In general, clusters can provide both high availability and scalability for business-critical applications. In today's market, experts estimate that approximately 80-90% of clusters are used to take advantage of the increase in server availability created by clusters.

The use of clusters to reduce computer system downtime bas a direct impact on a company's revenue and the MIS department budget. The impact depends on your calculated costs of downtime. In most cases, the cost of installing and maintaining a cluster are likely to be offset by the reduction in downtime costs, and, as businesses' reliance on computer systems intensifies, the cost of downtime is likely to increase.

Downloaded from www.Manualslib.com manuals search engine

ProLiant Cluster for NetWare 4.2 Overview

Novell products are designed to support incremental network growth, and the ProLiant Cluster for NetWare 4.2 continues this philosophy with its solid NetWare environment designed to provide both industrial-strength availability and scalability at the hardware, operating system, and application levels.

By deploying the ProLiant Cluster for NetWare 4.2, your company can reduce downtime and operating costs. At the same time, your network can benefit from the latest Internet/intranet and file and print services in a robust network environment for deploying general-purpose applications.

Features and Benefits of the ProLiant Cluster for NetWare 4.2

The ProLiant Cluster for NetWare 4.2 provides these features and benefits:

- Two node active-active cluster configuration
- Automatic failover and failback
- Protection against both hardware and software failures
- Failure detection over dedicated and/or non-dedicated links
- Support for any ODI™, IP compliant network interface card
- Support for Symmetric Multi-Processing (SMP)
- Industry-standard, non-proprietary hardware requirements
- High availability and lower total cost of ownership through:
  - Support for Fiber Channel Arbitrated Loop (FCAL)
  - Novell Directory Services (NDS) support
  - NetWare Loadable Module™ (NLM™)-based installation
  - Dynamic user licensing
  - Application fallback via IP

Novell's High-Reliability Solution

The ProLiant Cluster for NetWare 4.2, featuring Novell High Availability Server (NHAS), increases network reliability through the use of two NetWare 4.11 or NetWare 4.2 servers with a shared disk subsystem. By pairing two independent servers in an active-active configuration, the NHAS allows each server to access and mount the other server's data if there is a failure on a server in the cluster.

Each server monitors the status of the other server over three redundant lines of communication:

- NetWare Link – each server can monitor the presence of the other across the local network.
- Disk Link – NHAS uses quorum arbitration to allow each server to monitor the other server's shared storage connection and prevent concurrent access to the same data.
- Optional Dedicated Link – by creating a dedicated link between the two servers, each server can monitor the presence of the other more reliably, eliminating false failures caused by network traffic or prolonged disk inactivity.

Downloaded from www.Manualslib.com manuals search engine

The Failover Process

Costandby State

During normal operation, shown in Figure 1, each server node in the ProLiant Cluster for NetWare 4.2 operates as if it were a completely independent entity on the network. Server A and Server B are allocated their own dedicated volumes on the Compaq RAID Array 4000 shared storage. NHAS controls access for the shared storage and prevents one server from accessing the other server's volume. This controlling mechanism is vital to the operation of the cluster, since it prevents a *split-brain* scenario, wherein two servers may try to access the same data simultaneously; obviously, this could result in data corruption. In the given configuration, each server is operating in the *costandby* state, that is, independently but waiting for the other server to fail.

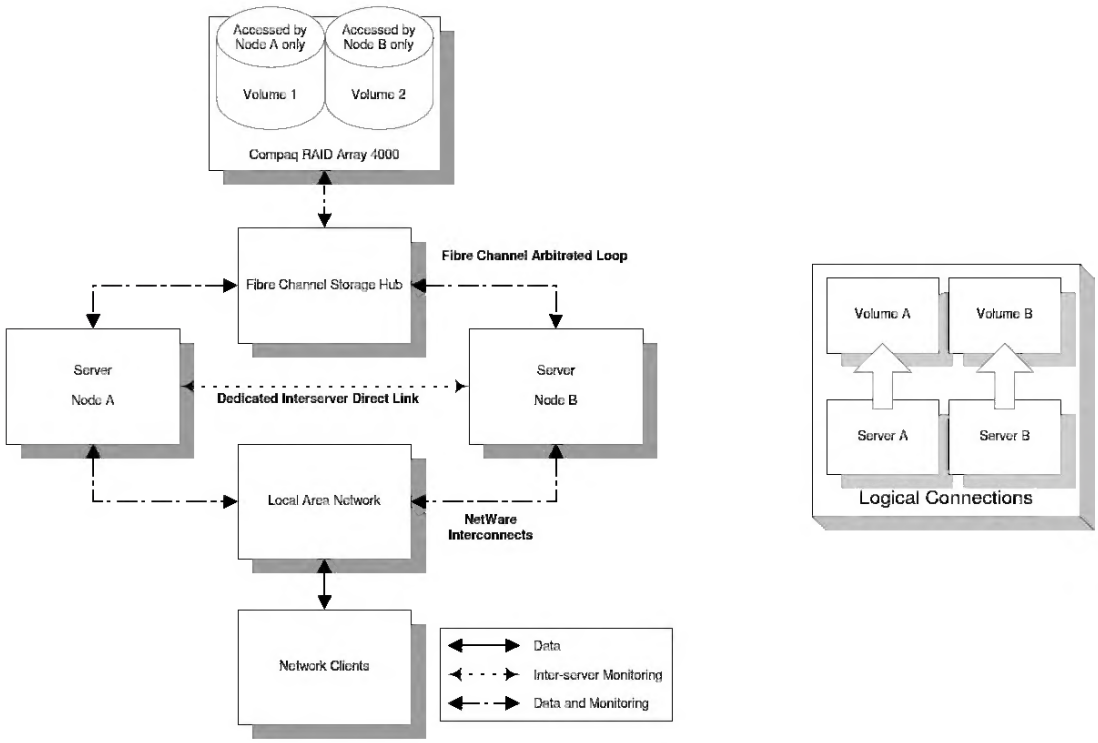

Figure 1. ProLiant Cluster for NetWare 4.2 Costandby State

Downloaded from www.Manualslib.com manuals search engine

Failover and Recovery States

Figure 2 illustrates the failover and recovery process. If no communication between the servers can be established across any of the three lines of communication, the server that has lost connections to the network and storage will enter *failover* state. At the same time, the other node, in this case, Server B, will enter the *recovery* state (Figure 2A). Upon entering the recovery state, Server B assumes responsibility for Server A volumes and clients (Figure 2B). Security and licensing are maintained by transferring NDS rights to the surviving server. Total failover time corresponds to the time necessary to detect the failure and complete the failover process. During a failover, the remaining server must mount the new volumes, launch any active server-based applications that were running on the failed device, and transfer security and license information. NHAS also allows manual failover of shared volumes at any time by pressing F at the NHAS console. By forcing a failover, you can remove a server from the network for scheduled maintenance without interrupting access to that server's volumes.

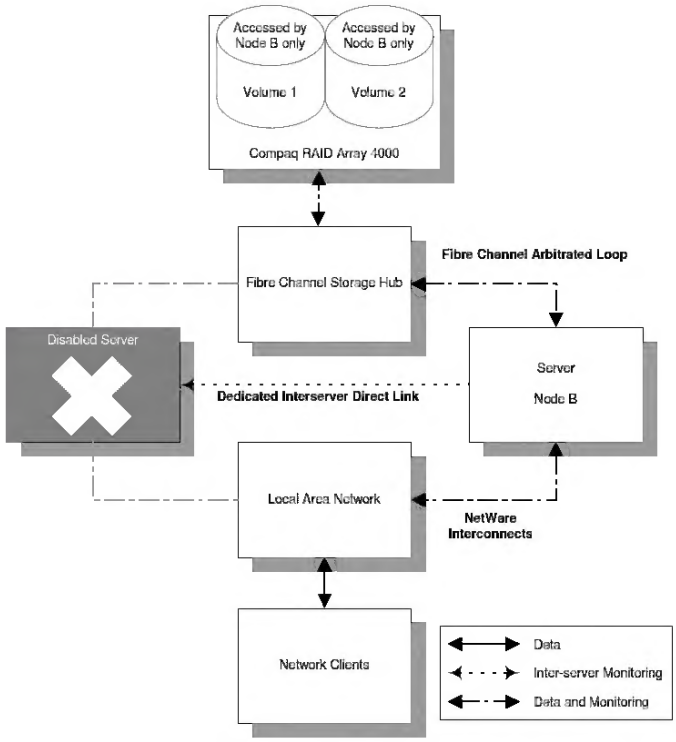

A. Node B Entering Recovery State

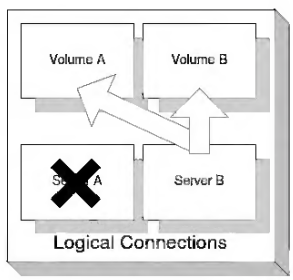

B. Node B Assuming Responsibility for Node A's Volumes

Figure 2. Failover and Recovery

Downloaded from www.Manualslib.com manuals search engine

Hardware Requirements

As shown in Figure 3, a typical ProLiant Cluster for NetWare 4.2 hardware configuration consists of two ProLiant servers with Fibre Channel Host Controllers, a Fibre Channel Storage Hub, and a Compaq RAID Array 4000 storage system with a Fibre Channel Array Controller.

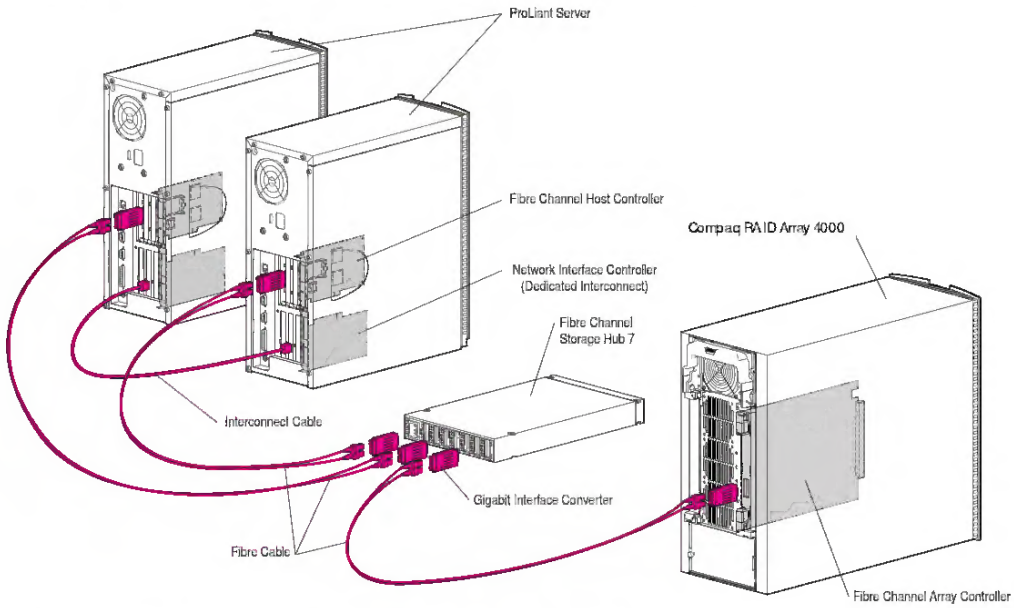

Figure 3. Typical ProLiant Cluster for NetWare 4.2 Configuration

Whether your intention is to upgrade an existing server or purchase a new Compaq server, the following minimum hardware requirements must be met for a ProLiant Cluster for NetWare 4.2 configuration.

Compaq ProLiant Servers

The primary components of a cluster are the servers. The initial release of the ProLiant Cluster for NetWare 4.2 supports a two-node cluster, where each node is a server. The key software component that imparts high availability to the cluster servers is Novell High Availability Server (NHAS). Throughout the development of the NHAS software, Compaq has been a partner with Novell to ensure that Compaq ProLiant servers are the ideal platforms for NHAS. Compaq has logged countless hours testing ProLiant clusters, which have successfully passed Novell's NHAS Certification test suite. This rigorous suite of tests ensures that the cluster works as an integrated unit, not just as a collection of individual components.

Downloaded from www.Manualslib.com manuals search engine

The following Compaq ProLiant servers bave been certified for use with both NHAS and Fibre Channel:

- Compaq ProLiant 1850R – Intel Pentium II
- Compaq ProLiant 3000 – Intel Pentium II
- Compaq ProLiant 5500R – Intel Pentium II Xeon
- Compaq ProLiant 6400R – Intel Pentium III Xeon
- Compaq ProLiant 6500 – Intel Pentium II Xeon
- Compaq ProLiant 7000 – Intel Pentium II Xeon

NOTE: Check the Compaq website at http://www.compaq.com to obtain an up-to-date list of cluster-certified servers. Each server in the cluster must have one local hard drive. In addition, Compaq recommends 256MB of RAM in each server.

In addition to the increased availability of clustering, Compaq ProLiant servers include many reliability features that provide a solid foundation for effective clustered server solutions. They offer excellent reliability through redundant systems, hot-pluggable components, and server health and monitoring. The high-availability features of ProLiant servers are a critical foundation of an effective ProLiant Cluster for NetWare 4.2 deployment and the cornerstone of mission-critical computing.

Compaq StorageWorks RAID Array 4000 Storage System

The ProLiant Cluster for NetWare 4.2 is based on a cluster architecture known as Shared Storage Clustering, in which clustered servers share access to a common set of hard drives. Novell High Availability Server requires all shared data to be stored in an external storage system to implement the necessary quorum arbitration.

The Compaq RAID Array 4000 storage system brings significant advantages to clustering, including:

- 100MB/s throughput
- Server-to-hub and hub-to-storage distances of up to 500 meters for short-wave and 10 kilometers for long-wave (per Fibre Channel link, 25 kilometers cumulative for all links)
- Increased connectivity and ease of use
- Reliable transmission of data in shared configurations
- Hot-pluggable drives
- Hot-pluggable power supply (in redundant power supply configurations)
- Manageability through Compaq Insight Manager

By using standard short-wave Fibre Channel cabling, the Compaq RAID Array 4000 can be placed up to 500 meters from the Fibre Channel Storage Hub, and the Fibre Channel Storage Hub can be placed up to 500 meters from the server. By using long-wave Fibre Channel cabling,

Downloaded from www.Manualslib.com manuals search engine

these distances can be increased to 10 kilometers (and a total of 25 kilometers on the Fibre Channel loop).

NOTE: The maximum distance between a single server and the Fibre Channel Storage Hub is 10 kilometers. A NetWare connection or dedicated direct link will need to connect each cluster node. Connecting the nodes over long distances will require a combined use of Ethernet network devices, such as hubs, switches, or repeaters.

The Compaq RAID Array 4000 storage system consists of four primary components:

- Fibre Channel Array
- Fibre Channel Storage Hub
- Fibre Channel Array Controller
- Fibre Channel Host Controller

When planning your Compaq RAID Array 4000 storage system, remember to include any necessary cabling. Also, the Fibre Channel Array Controller is not a separate component; rather, the array controller is an integral part of the Fibre Channel Array.

Compaq Fibre Channel Array

The Compaq Fibre Channel Array is the storage chassis that contains the disk drives, power supply, and Fibre Channel Array Controller. The Fibre Channel Array can hold twelve 1-inch or eight 1.6-inch Wide-Ultra SCSI-3 drives. The Compaq Fibre Channel Array comes in both a rack and a tower model.

Each ProLiant Cluster for NetWare 4.2 must have at least one Compaq RAID Array 4000 storage system as external shared storage. As many as five Compaq RA4000s can be attached to the Fibre Channel Storage Hub 7 and, with the Fibre Channel Storage Hub 12, up to ten Compaq RA4000s are supported. In both cases, two ports on the hubs are used by the server connections. With the Compaq RA4000 connected to both servers through the Fibre Channel Storage Hub, you can independently manage volumes on the servers in the cluster as needed by your application and environment.

Compaq Fibre Channel Storage Hub

The Fibre Channel Storage Hub creates the data path that allows one or more Compaq RAID Array 4000 storage systems to communicate with one or more ProLiant servers. The Compaq ProLiant Cluster for NetWare 4.2 requires at least one Fibre Channel Storage Hub; both servers and all shared storage for the cluster must connect to this hub.

As the names imply, the Fibre Channel Storage Hub 7 and the Fibre Channel Storage Hub 12 contain seven and twelve ports, respectively. As mentioned, using the hubs in a cluster environment will consume two ports for server connections, allowing five Compaq RAID Array 4000 storage systems on the Fibre Channel Storage Hub 7 and 10 arrays on the Fibre Channel Storage Hub 12.

For higher availability, two Fibre Channel Storage Hubs can be configured as a redundant pair. When used in combination with redundant Fibre Channel Array Controllers and redundant Fibre Channel Host Controllers, a complete backup data path can be configured between the cluster servers and the Compaq RAID Array 4000, as described below under "Redundant Fibre Channel Loop for Increased Availability."

Downloaded from www.Manualslib.com manuals search engine

NOTE: Currently, Fibre Channel Storage Hubs cannot be daisy chained, that is, hubs cannot be connected in series.

Compaq Fibre Channel Array Controller

One or two Fibre Channel Array Controllers, which reside in the Compaq RAID Array 4000 storage system, provide RAID and caching capabilities and provide the Fibre Channel interface to connect to the server. An important aspect of a higb-availability system is fault tolerance, traditionally accomplished by implementing RAID technology. Hardware RAID is an integrated part of the Compaq Fibre Channel Array Controller, located within the Fibre Channel Array chassis.

The Fibre Channel Array Controller supports RAID 0, 1, 4, and 5, as well as online capacity expansion and pre-failure notification for hard drives. In addition, the Fibre Channel Array Controller provides ECC (Error Checking and Correcting) memory for user-configurable read and write cache. The cache on the Array Accelerator is equipped with onboard rechargeable batteries, ensuring that the cached data is safe even with equipment failure or power outage. For a complete list of features and accompanying descriptions, refer to the Fibre Channel information on the Compaq website, http://www.compaq.com.

Typically, a Compaq RAID Array 4000 storage system is configured with a single controller. However, for a higher level of availability, two array controllers can be configured as a redundant pair, sucb that one controller is a backup for the other, as described under "Redundant Fibre Channel Loop for Increased Availability," below.

Compaq Fibre Channel Host Controller

The host controller, available as a PCI adapter, interfaces between the server and the Fibre Channel Storage System througb the fiber optic cable. At least two PCI Fibre Channel Host Controllers, one for each cluster node, are required in your Compaq ProLiant cluster. For higher availability, Fibre Channel Host Controllers can be configured in redundant pairs, sucb that one controller is a backup for the other, as described under "Redundant Fibre Channel Loop for Increased Availability," below.

Redundant Fibre Channel Loop for Increased Availability

For increased availability, a redundant Fibre Channel loop can be configured, which provides two redundant data paths between each cluster server and the Compaq RAID Array 4000 storage system. One data path is configured as the primary path and is initially active. The other data path is configured as a standby path and is initially inactive. The standby data path serves as a backup for the primary path. The redundant data path configuration requires the following additional components:

- A second, standby Fibre Channel Host Controller installed in each server
- A second, standby Fibre Channel Array Controller installed in the Compaq RA4000
- Recommended but not required: A second, standby storage hub with connections to the standby host controllers and to the standby array controller.

Figure 4 shows an example of this configuration.

Downloaded from www.Manualslib.com manuals search engine

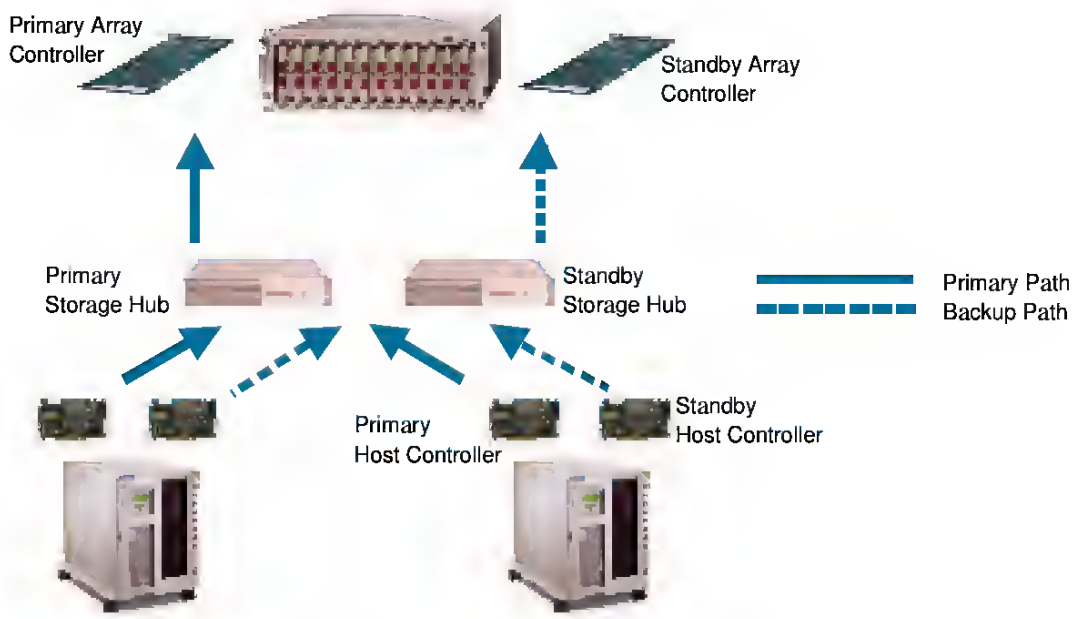

Figure 4. Redundant Fibre Channel Loop

As the figure shows, a redundant data path is configured between the cluster servers and the Compaq RAID Array 4000. The redundant components operate initially in standby mode. If any component (host controller, storage hub, array controller, or cable) in the primary path fails, the primary path becomes inactive, the standby path becomes active, and data transmission continues without interruption over the standby path.

Cluster Interconnects

The cluster interconnect is a data path over which nodes of a cluster communicate. This type of communication is termed *intra-cluster communication.* The ProLiant Cluster for NetWare 4.2 uses the standard public network controller card (NetWare Link), an optional Dedicated Link, and the Fibre Channel Disk Link to implement intra-cluster communication. At a minimum, the interconnect consists of two network adapters (one in each server) and a cable connecting the adapters to a public network hub to be utilized as the NetWare Link.

The cluster nodes use the interconnect data path to:

- Communicate cluster status
- Send and receive heartbeat

Since Novell High Availability Server does not mirror data between servers, the transfer of the heartbeat is very small (just the heartbeat or information about the health of the other server in the cluster is sent over the links).

IMPORTANT: NHAS 1.0 requires IPX as the cluster communication protocol for the Direct Link connection.

Downloaded from www.Manualslib.com manuals search engine

Network Adapters

Standard Compaq Ethernet adapters are the network adapters of choice for Compaq ProLiant clusters. Either 10Mbit/sec or 100Mbit/sec Ethernet may be used.

Software Components

Novell Software

Novell offers Novell High Availability Server to complement Novell NetWare 4.11 or NetWare 4.2. As of June 1999, NHAS will only run with Novell NetWare 4.11 and NetWare 4.2. Previous versions of NetWare are not supported.

Compaq Software

Compaq offers an extensive set of features and optional tools to support effective configuration and management of your clustered servers:

- Compaq SmartStart and Support Software CD
- Compaq System Configuration Utility
- Compaq Insight Manager
- Compaq Support Software Diskette for Novell NetWare (Novell SSD)

Compaq SmartStart and Support Software CD

SmartStart is the intelligent way to configure your Compaq ProLiant Server. SmartStart uses a step-by-step process to configure your server, load the system software, and integrate Compaq utilities, thereby ensuring that your server delivers maximum dependability and supportability. For information concerning SmartStart, refer to the Compaq Server Setup and Management pack.

Compaq System Configuration Utility

The SmartStart and Support Software CD also contains the Compaq System Configuration Utility. This utility is the primary means to configure hardware devices in your server, such as I/O addresses, boot order of disk controllers, etc. For information concerning the Compaq System Configuration Utility, see the Compaq Server Setup and Management pack.

Compaq Insight Manager

Compaq Insight Manager, loaded from the Compaq Management CD, is an easy-to-use software utility for collecting server information. Insight Manager performs the following functions:

- Monitors fault conditions and system status
- Monitors shared storage and interconnect adapters
- Forwards server alert fault conditions
- Allows remote server administration

In Compaq servers, each hardware subsystem, such as disk storage, system memory, and system processor, has a robust set of management capabilities. Compaq Full Spectrum Fault

Downloaded from www.Manualslib.com manuals search engine

Management notifies of impending fault conditions and keeps the server up and running in the unlikely event of a hardware failure. For information concerning Compaq Insight Manager, see the Compaq Server Setup and Management pack.

Compaq Support Software Diskettes for Novell NetWare (Novell SSD)

The Compaq SSD for Novell NetWare contains device drivers and utilities that enable you to take advantage of specific capabilities offered on Compaq products. The drivers and utilities are provided for use only with Compaq hardware.

For additional details on the features contained in the Novell SSD, visit

http://www.compaq.com/support/files/server/softpaqs/Netware/NSSD.html

The Compaq SSD for Novell NetWare is included in the Compaq Server Setup and Management pack.

Downloaded from www.Manualslib.com manuals search engine

Planning Novell High Availability Server Installation

Listed below are some guidelines to keep in mind before installing NHAS on the ProLiant Cluster for NetWare 4.2:

- It is imperative that the volumes created in the shared storage are not mounted on both servers. Depending on your requirements, you may want to allocate a volume to each server, keeping in mind that NHAS will later prohibit the other server from accessing that volume until a failover occurs.

  To check the status of each server's allocated storage device, at the command prompt issue the following commands in the order listed:

  CONFIG
  SCAN FOR NEW DEVICES
  LIST DEVICES
  VOLUME

- Be sure to install the latest NetWare Support Pack, which is available from Novell, and the Compaq Support Software for Novell NetWare (Novell SSD).
- The installation of NHAS requires administrative privileges.
- The servers in the cluster must reside in the same NDS tree to fully utilize the NDS capabilities. Note that NHAS does not require that the servers in the cluster reside in the same NDS container, but it is recommended for manageability sake.
- Each server in the cluster must have an internal hard drive that is not part of the shared subsystem. If NetWare is installed with the shared drive subsystem connected, the installer should make sure that SYS: volumes of both clustered servers are installed on the internal, non-shared drives. Having the SYS: volumes on the shared subsystem would be confusing and potentially damaging.
- NetWare volumes can span multiple drives, but if a volume spans more than one drive, all of the drives that it spans must be shared, or all of them must be not shared. It is also possible to have more than one volume on a single shared NetWare partition. Some hardware configurations will allow volumes to be mounted on multiple servers at the same time but once NHAS is installed this is not allowed. If a single volume is mounted on multiple servers, at the same time, data loss will occur.

Verification of NHAS Installation

These procedures will help you validate a proper NHAS installation on both NHAS servers.

- After the NHAS installation, restart your Compaq ProLiant server. NHAS should load, and you should have the capability to switch to the HA Server screen by pressing **Ctrl+Esc** and selecting the HAServer Screen from the menu
- At the HA Server screen, there will be generalized information about the status of your NHAS cluster. Make sure that the SHARED VOLUME is MOUNTED, and the SERVER is in COSTANDBY mode on both servers. Also, verify that NETWARE LINK, DISK LINK, and DIRECT LINK (if utilized) are all labeled CONNECTED.

Downloaded from www.Manualslib.com manuals search engine

- To perform a simple verification of operation, press **F** on the keyboard. Compaq recommends that you pause the NHAS software during these tests, by typing **P**, so an automatic failback cannot take place. Pressing **F** will force a server to fail. The failed server should report the status as FAILOVER. Soon after the induced failure, both servers should recognize that they are both operational. The surviving server will return to COSTANDBY state when the previously failed server enters the RECOVERY state. If successful, both servers in the NHAS cluster will fallback to COSTANDBY state and all quorum arbitration links will be connected.
- Some other commands to note:

  C - Changes the first server into the COSTANDBY state.

  ESC - Used to exit the program. When the Escape button is pressed, you are presented with a confirmation prompt. Choose Yes to exit, or choose No to remain in the application.

  F - Initiates a manual failover on the first server. When you press f, you are prompted to confirm that you want to manually failover the first server and acquire the disk devices from the second server.

  H - Displays the Help screen.

  P - Pauses the AutoSwitch feature. When you press **P**, the AutoSwitch feature is paused and automatic failover or rollback will not take place. When you pause NHAS, the status turns red to indicate that the server cannot change states. To unpause the NHAS Console, press p again. Note that the Pause feature does not affect a transitional state. For example, once the failover process has reached the ENTERING FAILOVER state, the Pause feature will not stop the failover.

  R - Cbanges the server into the RECOVERY state without changing the second server to the FAILOVER state. Any shared volumes are dismounted and any shared drives are released. If you want to failover the system, press **F** on the second server, which will force first server into the RECOVERY state.
- If the failover was unsuccessful, check the network configuration and verify that it is properly set up.
- If further problems persist, check the *AUTOEXEC.NCF* file for the following entries:

  ```
  REM HA Server Installation
  Search Add Sys:\HASERVER
  Set Auto Restart After Abend=0
  Load VincaCom
  Load VinclPX (master/slave)=address, card SLUNBIND
  Load HASERVER
  LOAD HASTRUST
  ```

  If any of these entries are incorrect, you should either edit the file to reflect the correct information or reinstall NHAS.
- If problems continue, check SYS:\HASERVER*HASERVER.LOG* for errors. Contact your local support representative if necessary.

Downloaded from www.Manualslib.com manuals search engine

- If you wish to cause a manual failover on the cluster containing Server A and Server B, with Server B being the surviving server, initiate the failure from the console of Server A. Server B would then be the surviving server and Server A would enter recovery.
- If the Disk Link fails, then there is a more serious problem somewhere along the shared drive channel. If the problem is in the channel between one server in the cluster and the shared drive subsystem, then that server will be disarmed and its shared drives will have already become inaccessible. The volumes will not be mounted on the other server unless all other links fail or a manual failover is invoked. If the problem is in the shared subsystem, then some or all of the shared drives will become inaccessible and one or both servers will be disarmed. A Disk Link failure alone will not cause a server failover.

Downloaded from www.Manualslib.com manuals search engine

ProLiant Cluster for NetWare 4.2 Administration

Outlined below are some helpful guidelines to assist you in administering and supporting the ProLiant Cluster for NetWare 4.2.

Increasing Availability

Intra-cluster Communication

Should the servers experience a disruption in intra-cluster communication, they will initiate the failover process. Obviously, this condition is only acceptable in the event of an actual failure that disrupts the server's ability to service client requests. Therefore, the reliability of the connections between the servers in the cluster is a direct measure of the clustered system's ability to maintain costandby operation, wherein both servers are fully operational.

There are several ways to reduce the possibility of a disruption of intra-cluster communication. Each method shares an underlying objective: to create redundant paths over which monitoring will continue if the another path is disrupted. The ProLiant Cluster for NetWare 4.2 provides this redundancy through the three lines of communication between servers: the NetWare link, the optional dedicated link, and the disk link.

Compaq extends this redundancy with a feature that configures two Compaq Ethernet adapters or two ports on a single adapter, such that one is a hot backup for the other. This feature, called Advanced Network Fault Detection and Correction, is available on all Compaq 10/100 Fast Ethernet products. This feature reduces the possibility of downtime due to network failure by allowing the administrator to configure a redundant interconnect for each communication path. In a redundant interconnect configuration, each communication path is configured with a primary link and a backup link. If a component in the primary link fails, communication continues over the backup link. The redundant interconnect feature significantly enhances the availability of clustered systems.

Quorum Arbitration

Each server uses quorum arbitration to monitor and verify that the other server is operational. By monitoring the other server's usage of a reserved region on the shared volumes, both servers can be sure that the other is still accessing the storage array. In the event that a server does not update the information in this shared region, the cluster becomes aware of a potential failure. In order to prevent false failovers, NHAS will verify the disk link status 4 times. If after the fourth verification, or roughly 20 seconds, Server B has not accessed the reserved region of the disk, Server A will set the disk link status to disconnected. If the NetWare link (and dedicated link, if implemented) status is also disconnected, NHAS will initiate a failover to Server A.

IP and Applications Failover

Configuration

If you want your IP-based applications to failover to the other server when their host server fails, you must set up your ProLiant Cluster for NetWare 4.2 to redirect all IP traffic from the failed

Downloaded from www.Manualslib.com manuals search engine

address to the surviving server. The following example provides basic information on the necessary steps to setup a simple IP failover scenario.

There are three key commands to add a failed IP address to a surviving server. They are ADD, DEL, and DISPLAY. Their formats and syntax are shown below:

```
add secondary ipaddress <failed server IP address>
del secondary ipaddress <failed server IP address>
display secondary ipaddress
```

In the examples above, *<failed server IP address>* is the IP address of the IP-based application to be failed-over if the host server fails. The ADD command assigns the IP address to the surviving server in the event of a failover. The DEL command removes the IP address from the server when recovery occurs. The DISPLAY command is for informational purposes only and can be used from the command line to view all secondary IP addresses.

The following lines below show the usage of the commands in a .NCF batch file one can configure. Comments are remarked with a "#" sign:

```
# enables duplicate IP addresses support under NetWare
set allow ip address duplicates=on
# adds all secondary IP addresses
add secondary ipaddress 123.45.67.8
add secondary ipaddress 123.45.67.9
# launches your application
run-now.nlm
# deletes all secondary IP addresses except for the one you want to launch
del secondary ipaddress 123.45.67.8
```

These .NCF batch files should be created in the search directory of each server node, such as SYS:\SYSTEM. Each server will have varying IP addresses input in their .NCF files to discern which secondary IP addresses are to be deleted and which are to be kept during a failover and recovery. IP handling during a failover should add the secondary IP to the surviving server so that it may take over client requests for the failed node. At the same time, NHAS IP handling should remove the same secondary IP from the failed node. Two separate batch files will be made for each server: one for the failover to a node and the other to failback to the recovered node. In addition, the *HASERVER.INI* file must contain the *FailOverCommand*= and *RecoveryCommand*= entries to automatically reconnect IP-based applications after failover. The names of the respective failover and recovery .NCF files must be added to the commands for a successful failover and recovery of the IP-based application.

Note: The sample *HASERVER.INI* on page 43 of the NHAS User's Guide is incorrect. See Novell's Errata sheet for the correct layout. This is only a document change. Your *HASERVER.INI* file should be in the correct state.

HASERVER.INI Information

To enable IP application failover, it is vital you know and understand the command lines inside of the *HASERVER.INI*, so that you can tune NHAS specifically to your requirements. Below is a sample *HASERVER.INI* with a description of each line:

Downloaded from www.Manualslib.com manuals search engine

[HAServer]	
Server=PL7000-2	This identifies the NetWare Servers in the NHAS Cluster.
Server=PL7000-1	This identifies the NetWare Servers in the NHAS Cluster.
NWLinkTimeout=30	This sets the NWLink timeout value before NHAS triggers the link down. Keep in mind that this value is also used when the disk and direct links fail.
FailoverDelay=30	This specifies the delay time of NHAS after the links are down to failover the disk volumes to the surviving cluster node.
Broadcast=Off.	Enabling this will broadcast a network message informing users the status of the servers in the cluster when a failover or recovery happens.
Rollback=Automatic	Setting this value to **manual** will disable automatic failback of the volumes should the server that went down go back online.
DiskLink=258E5776	This is a unique number used for quorum arbitration by the disk link. Note that if this value is accidentally changed or erased, we recommend that you reinstall NHAS.
[Server.PL7000-2]	This subsection details the information for this particular server, **PL7000-2**.
Volume=VOL2	Volume settings identify which volumes should remain exclusive to this server.
Volume=SYS	Volume settings identify which volumes should remain exclusive to this server.
PrimaryLicense=20614660	This License Number is used as the main Server License for this node during normal operating conditions.
RecoveryCommand=REC1.NCF	By specifying a command or an executable script, when a server enters the RECOVERY STATE, i.e. has failed, it processes the commands listed here. This specified command should call a file that contains all commands necessary for a successful IP failover.
FailoverCommand=FAIL1.NCF	By specifying a command or an executable script, when a server enters the FAILOVER STATE, i.e. is taking over for the other failed server, it processes the commands listed here. This specified command should call a file that contains all commands necessary for a successful IP failover.
CostandbyCommand=CO2.NCF	By specifying a command or an executable script, when a server enters the COSTANDBY STATE, i.e. Normal condition, it processes the commands inputted here. This specified command should call a file that contains all commands necessary for a successful IP failover.
DiskLinkSector=0	This specifies the segment of the disk used for quorum arbitration by the Disk Link.
SecondaryLicense=60022796	This License Number is used to append the other Server's License should a failover happen.
[Volume.PL7000-2.VOL2]	This subsection details the information for this particular server, **PL7000-2**, in respect to Volume 2.
Shared=True	Setting the value to **false** will disable the ability for a drive to failover.
Mounted=True	Setting the value to **false** will disable NHAS from automatically mounting this volume.

Downloaded from www.Manualslib.com manuals search engine

HotfixIdentifier=258E5776	This value is derived from the partition time stamp NetWare creates. It is used to recognize volumes for NHAS.
MountCommand=	By specifying a command or an executable script, when a server mounts this particular volume, it processes the commands listed here.
DismountCommand=	By specifying a command or an executable script, when a server dismounts this particular volume, it processes the commands listed here.
[Volume.PL7000-2.SYS]	This subsection details the information for this particular server, **PL7000-2**, in respect to the SYS: volume.
Shared=False	This value should stay false.
Mounted=True	Setting the value to **false** will disable NHAS from automatically mounting this volume.
HotfixIdentifier=258E5AE8	This value derive from the partition time stamp NetWare creates and is used to recognize volumes for NHAS.
MountCommand=	By specifying a command or an executable script, when a server mounts this particular volume, it processes the commands inputted here.
DismountCommand=	By specifying a command or an executable script, when a server dismounts this particular volume, it processes the commands inputted here.
[Server.PL7000-1]	This subsection details the information for this particular server, **PL7000-1.**
Volume=VOL1	These Volume settings identify what should remain exclusive to this server.
Volume=SYS	These Volume settings identify what should remain exclusive to this server.
PrimaryLicense=60022796	This License Number is used to as the main Server License for this node during normal operating conditions.
RecoveryCommand=REC1.NCF	By specifying a command or an executable script, when a server enters into the RECOVERY STATE, i.e. has failed, it processes the commands inputted here. This specified command should call a file that contains all commands necessary for a successful IP failover.
FailoverCommand=FAIL1.NCF	By specifying a command or an executable script, when a server enters into the FAILOVER STATE, i.e. is taking over for the other failed server, it processes the commands inputted here. This specified command should call a file that contains all commands necessary for a successful IP failover.
CostandbyCommand=CO1.NCF	By specifying a command or an executable script, when a server enters into the COSTANDBY STATE, i.e. Normal condition, it processes the commands inputted here. This specified command should call a file that contains all commands necessary for a successful IP failover.
DiskLinkSector=1	This specifies the segment of the disk used for quorum arbitration by the Disk Link.
SecondaryLicense=20614660	This License Number is used to append the other Server's License should a failover happen.

Downloaded from www.Manualslib.com manuals search engine

[Volume.PL7000-1.VOL1]	This subsection details the information for this particular server, **PL7000-1**, in respect to Volume 1.
Shared=True	In order for volumes to failover properly, this must be set to **true** Setting the value to **false** will disable the ability for drive to failover.
Mounted=True	Setting the value to **false** will disable NHAS from automatically mounting this volume.
HotfixIdentifier=258E576C	This value, derived from the partition time stamp NetWare creates, is used to recognize volumes for NHAS.
MountCommand=	By specifying a command or an executable script, when a server mounts this particular volume, it processes the commands inputted here.
DismountCommand=	By specifying a command or an executable script, when a server dismounts this particular volume, it processes the commands inputted here.

[Volume.PL7000-1.SYS]	This subsection details the information for this particular server, **PL7000-1**, in respect to the SYS: volume.
Shared=False	This value should stay false.
Mounted=True	Setting the value to **false** will disable NHAS from automatically mounting this volume.
HotfixIdentifier=258E5780	This value, derived from the partition time stamp NetWare creates, is used to recognize volumes for NHAS.
MountCommand=	By specifying a command or an executable script, when a server dismounts this particular volume, it processes the commands inputted here.
DismountCommand=	By specifying a command or an executable script, when a server dismounts this particular volume, it processes the commands inputted here.

Sample IP Failover

As the clearest method of explaining the IP failover procedure, the following sample is a walkthrough of a failover performed during a client ping. Hopefully, this example will shed light on the relationship between the *HASERVER.INI* parameters and the operation of the cluster.

The following example describes a series of files that explain how to implement IP failover between the two servers in the cluster. The script files that follow should be used as a reference when creating your own custom IP failover script files.

Scenario

IP address of Server A (PL7000-1): 123.45.67.1 with VOL1: on shared disk

IP address of Server B (PL7000-2): 123.45.67.2 with VOL2: on shared disk

Virtual IP address of VOL1: 123.45.67.11

Virtual IP address of VOL2: 123.45.67.22

Client IP address: 123.45.67.5

Downloaded from www.Manualslib.com manuals search engine

Configure both servers for a seamless IP failover.

1. Create a file called *CO1.NCF* on Server A with the following entries (the lines beginning with # are comments for explanation and clarity):

```
# deletes the other server's secondary IP addresses
del secondary ipaddress 123.45.67.22
# add all secondary IP addresses
add secondary ipaddress 123.45.67.11
# display all secondary IP addresses
display secondary ipaddress
```

   Then on Server A, edit the file *HASERVER.INI* and modify the CostandbyCommand parameter, under the section [Server.PL7000-1], as shown below:

```
CostandbyCommand=CO1.NCF
```

   Save this file, *HASERVER.INI*, in the SYS:\SYSTEM directory. Whenever a server goes into the costandby state, this script will configure the server to respond to all client requests to 123.45.67.11 and, at the same time, removing any secondary IP addresses for the other server.

2. Create a file called *CO2.NCF* on Server B with the following entries (the lines beginning with # are comments for explanation and clarity):

```
# deletes the other server's secondary IP addresses
del secondary ipaddress 123.45.67.11
# adds all secondary IP addresses
add secondary ipaddress 123.45.67.22
# displays all secondary IP addresses
display secondary ipaddress
```

   Then on Server B, edit the file *HASERVER.INI* and modify the CostandbyCommand parameter, under the section [Server.PL7000-2], as shown below:

```
CostandbyCommand=CO2.NCF
```

   Save this file, *HASERVER.INI*, in the SYS:\SYSTEM directory. Whenever a server mounts goes into the costandby state, this script will configure the server to respond to all client requests to 123.45.67.22 and, at the same time, removing any secondary IP addresses for the other server.

3. Create a file called *FAIL1.NCF* on Server A with the following entries (each command is preceded by remarked comments):

```
# adds all secondary IP addresses
add secondary ipaddress 123.45.67.22
# displays all secondary IP addresses
display secondary ipaddress
```

   Then on Server A, edit the file *HASERVER.INI* and modify the FailoverCommand parameter, under the section [Server.PL7000-1], as shown below:

```
FailoverCommand=FAIL1.NCF
```


Downloaded from www.Manualslib.com manuals search engine

Save this file, *HASERVER.INI*, in the SYS:\SYSTEM directory. Whenever Server B enters the recovery state, this script will have NHAS configure Server A to respond to all client requests to 123.45.67.22.

4. Create a file called *FAIL2.NCF* on Server B with the following entries (each command is preceded by remarked comments):

```
# adds all secondary IP addresses
add secondary ipaddress 123.45.67.11
# displays all secondary IP addresses
display secondary ipaddress
```

   Then on Server B, edit the file *HASERVER.INI* and modify the FailoverCommand parameter, under the section [Server.PL7000-2], as shown below:

```
FailoverCommand=FAIL2.NCF
```

   Save this file, *HASERVER.INI*, in the SYS:\SYSTEM directory. Whenever Server B enters the recovery state, this script will have NHAS configure Server B to respond to all client requests to 123.45.67.11.

5. Create a file called *REC1.NCF* on Server B with the following entries (each command is preceded by remarked comments):

```
# adds all secondary IP addresses
del secondary ipaddress 123.45.67.11
# displays all secondary IP addresses
display secondary ipaddress
```

   Then on Server B, edit the file *HASERVER.INI* and modify the RecoveryCommand parameter, under the section [Server.PL7000-2], as shown below:

```
RecoveryCommand=REC1.NCF
```

   Save this file, *HASERVER.INI*, in the SYS:\SYSTEM directory. Whenever Server A enters the recovery state, this script will have NHAS configure Server A to no longer respond to client requests to 123.45.67.22.

6. Create a file called *REC2.NCF* on Server A with the following entries (each command is preceded by remarked comments):

```
# adds all secondary IP addresses
del secondary ipaddress 123.45.67.22
# displays all secondary IP addresses
display secondary ipaddress
```

   Then on Server A, edit the file *HASERVER.INI* and modify the RecoveryCommand parameter, under the section [Server.PL7000-1], as shown below:

```
RecoveryCommand=REC2.NCF
```

   Save this file, *HASERVER.INI*, in the SYS:\SYSTEM directory. Whenever Server B enters the recovery state, this script will have NHAS configure Server B to no longer respond to client requests to 123.45.67.11.

7. Ping Server A with the virtual IP 123.45.67.11 from a client.

   Accordingly, the client ping should receive a reply from the server (volume).

Downloaded from www.Manualslib.com manuals search engine

8. Initiate a failover to Server B in the NHAS console. In other words, we will cause Server A to fail and transfer all network services to Server B. For the purposes of this illustration, cause Server A to fail by powering off the server. Initiating a failover by pressing **F** at the console requires that the NHAS software be paused on the downed server to prevent immediate recovery. Powering the server off ensures that the failover will be successful and sustained.

   At this point, Server B will execute the *FAIL2.NCF* file which will assign Server A's former secondary IP address as a secondary IP address handled by Server B.

   Throughout the process, the client ping should be able to reach 123.45.67.11. After the nearly instantaneous transfer is complete, Server B alone will respond to any requests sent to both 123.45.67.11 and 123.45.67.22

   In essence, the client is now talking to Server B exclusively. If a failback to Server A is initiated at the console, Server B executes the *CO2.NCF* script file, which deletes the 123.45.67.11 secondary IP transferred from Server A. At the same time, Server A will process the *CO1.NCF* file and, once again, handle requests for IP address 123.45.67.11. With the failback complete, Server B will handle only requests to 123.45.67.22.

Notes on Failover

As an emerging technology, clustering presents unique challenges to legacy applications and environments. While none of these challenges is prohibitive, certain steps or precautions are often necessary to accommodate the failover process.

Server Applications and the SYS: Volume

Some applications, such as Netscape Fast Track Server, write the executable files to the server's SYS: volume. In the event of a failover, the surviving server would no longer be able to access the files on the downed server's SYS: volume. If your ProLiant cluster runs applications that write to the server's SYS: volume, you must install the executables for that application on both servers. In addition, you must create a script that will launch any such applications on the surviving server after a failover.

Client Drive Mappings with NDS and Novell Client 32

If Novell Client32 and NDS are being used, the clients will remain attached to the network during a cluster failover. Network drives that were mapped attached to drives previously owned by the failed server, though, will be broken. The existing Novell Client32 technology does not support automatic failover drive mappings. A failed-over shared drive can be re-mapped in Windows by using Windows Explorer as follows:

1. Right click on the Start button.
2. Select Explore.
3. Select Map Network Drive from the Tools menu. Then select a drive letter and specify the path to the data to be mounted as a drive on the client.
4. Finally, choose whether you want to reconnect this drive mapping. Then select OK.

   If the client applications are IP only, and the virtual IP addresses are transferred at failover to the surviving server, then these applications will reconnect after the failover is complete.

Downloaded from www.Manualslib.com manuals search engine

Z.E.N.works Application Failover

Any server applications accessing the shared drives after a failover must be started on the surviving server as described above, and any network drives will need to be re-mapped. It is possible using Novell Z.E.N.works with the Novell Application Launcher to set up an application in NDS with failover capabilities. Using Z.E.N.works, a network administrator can create an application object in NDS for each server and then allow the failover of the application in the NetWare Administrator to point to the other server's NDS object. At failover, the clients' application will disconnect or crash if it was using resources on the failed server, but the user will be able restart the application from the failed over server. The only data loss would be what was not saved before the application disconnects or crashes. IP only clients and applications will be disrupted during the failover process, but if the IP configuration information and application are transferred to the surviving server, the client will resume normal operation after failover.

Disruptions During Failover

With NHAS, the recovery time in the event of a failover depends on the types of applications running on the server. Client/server applications performing database transactions in which records are being accessed by the database, or word processing applications where the server is accessed only to read and write a document, may not see any problem during a server failure and recovery.

Upon failover, some disruption to applications can occur. This may vary from a slight delay in accessing the recovery server to the need for complete re-logging into the system. Other applications, for example, database servers, must have clustering awareness built into them so that transactions can be rolled back and logs parsed to ensure data integrity is maintained. Oracle, on the other hand, features complex transaction logging systems that maintain the integrity of the database in the event of a failure, regardless of cluster awareness. But these interruptions are indeed trivial compared to the delay and cost of manually swapping in a new server and restoring the data and applications to the last backup.

User Licenses in a Failover

Novell High Availability Server's support for dynamic user licensing automatically adds the user license count from a failed server to your surviving server at failover time. When the failed server is brought back into operation, the user license count is automatically restored.

Management

Failover

Once a server initiates the failover detection process, a failover can only be paused. During the failover countdown, you may press **P** to pause the countdown. However, once unpaused and having expired the failover delay, the server begins "Entering Failover." At this point, failover cannot be stopped. If Server A does not generate any traffic across any line of server-to-server communication because of an intensive thread or other condition, it is possible for Server B to detect that Server A has failed, if the period of traffic inactivity continues. This "partition in time" may actually initiate a failover and cause all volumes of Server A to be dismounted and then mounted as volumes on Server B. However, Server A has not truly failed. Moreover, since Server A has been too busy with this internal thread, it has not noticed that Server B has assumed control of its resources.

Downloaded from www.Manualslib.com manuals search engine

This state could be very damaging because, at this point, both servers have assumed control of the same shared drives. NHAS handles this situation by immediately recognizing this scenario and updating both servers to the true state of the cluster. As a result, as soon as Server A has resumed network communication, NHAS forces Server A into recovery mode and dismounts Server A's volumes. When the failover process has completed, NHAS will initiate the restoration of Server A and place both servers in costandby mode.

Compaq recommends that system integrators/administrators test their clustered configuration for this and other conditions to ensure proper configuration. You may cbeck the NHAS log file for console screen messages to review the response of your cluster to any tests. The console messages are logged in the *HASERVER.LOG*, which can be found in the SYS:\HASERVER directory.

NDS

It is not necessary for servers in the cluster to contain NDS replicas as long as trusted replicas of this NDS tree exist on the network. If one server in the cluster contains a trusted NDS replica then it is recommended, for availability reasons, that both servers contain a replica. If the two servers in the cluster are the only servers in the NDS tree and one contains the master record of the NDS tree, the other server should contain an NDS replica to preserve high availability of NDS. It doesn't matter whether the servers in the cluster reside in the same NDS organizational unit, but its is recommended that both servers of the ProLiant cluster reside in the same tree. NHAS translates the directory and volume trustee rights for the files and directories on the volume as per the NDS tree. This is performed by the *HASTRUST.NLM*.

Should Server A fail and its shared drive failover to Server B, the newly mounted volumes must be manually added using one of the Novell NDS administration utilities to reflect the changes. Using NWAdmin is the recommended method. Long name spaces, Mac, NFS, etc. on a shared volume will transfer during a failover, but you need to make sure that the name space support modules for all name spaces in the cluster are loaded on all servers in the cluster. Name space support modules have a .NAM file extension, for example, *LONG.NAM*.

When a server fails, any volume or directory information written on the drives at the time of failure might be damaged and in need of repair before it can be mounted by the surviving server. To protect against possible down time, you must bave NetWare automatically run Vrepair on damaged volumes. NetWare performs this task by default unless you have disabled this feature. To check the status of this feature, enter *set automatically repair bad volumes* at the server console. If enabled, ON is returned. If OFF is returned, you can follow the instructions on the screen to enable this option. Additionally, when a server fails, a volume may be left in an undetermined state known as *out of sync*. In order to resynchronize a volume, you should use the NetWare installation utility.

Print Services

NHAS is an active/active data failover environment so both servers can actively be used for file, print, or application services. Unfortunately, with NetWare 4.11/IntraNetWare and NetWare 4.2, Novell legacy printing will function with NHAS but print queues assigned to shared volumes will not transfer in a failover. Novell did allow print queues to be assigned to specific volumes, but with their legacy printing product, the print queues are still internally assigned to a server name and volume. With the server name tied to the print queue it is not possible to transfer queued print jobs to the surviving server. It is strongly recommended that NHAS customers use Novell Distributed Print Services (NDPS) which fully supports cluster failovers.

Downloaded from www.Manualslib.com manuals search engine

Tuning NHAS

Memory Requirement Calculation

Since NHAS runs on NetWare 4.11 and NetWare 4.2, most of the performance tuning for a cluster is handled by tuning the OS itself. However, the fact that NHAS has two modes of operation adds a twist to system requirements. Specifically, the optimal amount of RAM that should be installed in each clustered server will vary drastically between costandby and failover modes. As with NetWare itself, the recommended amount of RAM depends on a number of factors: size of the internal drive(s), size of all shared volumes, the applications and name spaces loaded, number of users, etc. A good reference article on server memory calculation can be found on the Novell web site at:

http://developer.novell.com/research/appnotes/1995/November/03/index.htm

There is also a DOS based server memory calculator available through the Novell web site called *SMEM.EXE*. Whichever method you choose to calculate RAM requirements, remember that you should calculate using a failover condition, where all shared volumes are mounted on one server, as the basis.

Caching Concerns

As a matter related to performance, NetWare uses a file cache to increase server performance. However, when this cache is used and a server fails, information that is contained in the file cache and not yet written to disk is lost. Configure the NetWare file cache to balance your performance needs against the possible loss of critical data.

For more information on tuning NHAS to your system requirements, see the documentation accompanying the software.

Technical FAQs

What happens to a client running server based applications when a server failover occurs?

NetWare High Availability Server can use scripts or the NetWare Application Launcher, a Z.E.N.works component, to restart any server-based applications that may have been shut down when the server failed. These applications will restart on the surviving server, allowing the client to continue working from the other server node. Most client-based applications will continue to function normally, with a slight pause while the volume is reconnected through the persistent links. Some applications will require reloading to continue operation.

What happens when the client is running local applications to work on a document located on the server node that failed?

Most applications will continue to function normally. Clients with a *stateful* connection, such as IPX, will not be able to reconnect the volume that was previously in use since the volume is referenced differently in the NDS tree. The client will not have to reconnect, but will have to select the new volume object under the other server node in the NDS tree.

Downloaded from www.Manualslib.com manuals search engine

Can data be lost in documents that have not been saved frequently if the server supporting the volume fails?

Some client applications may fail when the connection to the document is lost. Even though this data may still be in the client's RAM, data loss could result.

Will the behavior differ for different clients?

Yes. The behavior will differ between clients depending on the type of connection; *stateful* vs. *stateless.* As previously stated, *stateful* connections must be reestablished. However, *stateless* connections, like TCP/IP database clients will seamlessly reconnect to the clustered service. Both types of connections can use the NetWare Application Launcher to restart any server-based applications needed by the failed-over clients.

Where can one find info on these products?

http://www.novell.com/products/clusters/

and

http://www.compaq.com/solutions/enterprise/highavailability.html

What happens to printer jobs when a system running Novell High Availability Server fails?

Any print jobs still in RAM will be lost. The other server will resume all others that have been written to disk when the volume is mounted on the other active server.

What is the "slunbind" command line parameter on the VINCAIPX line of the AUTOEXEC.NCF?

With StandbyServer ™ only one server is monitoring the other, so monitoring communication is unidirectional. NHAS uses the same monitoring technology but requires bi-directional communication because each server in the cluster is monitoring the other. The slunbind command enables both servers in the cluster to function properly during a failure.

How do I uninstall NHAS?

The first release of NHAS does not include an automated uninstall. To uninstall NHAS, the administrator must delete the SYS:\HASERVER directory, then delete the NHAS command lines from the AUTOEXEC.NCF.

How many times must the Validate NLM be run to validate the cluster?

NHAS must be validated within 30 days of being installed through the Validate NLM and a Validation Key that is specific to each NetWare serial number. Validate NLM must be run once on each server, for a total of two times.

Downloaded from www.Manualslib.com manuals search engine